arXiv:1611.08198v1 [cs.DS] 24 Nov 2016

Burrows-Wheeler transform and LCP array construction in constant space[☆]

Felipe A. Louza[a,*], Travis Gagie[b], Guilherme P. Telles[a]

[a] *Institute of Computing, University of Campinas, Campinas, SP, Brazil*
[b] *School of Informatics and Telecommunications, Diego Portales University, Santiago, Chile*
CeBiB (Center of Biotechnology and Bioengineering), Santiago, Chile

Abstract

In this article we extend the elegant in-place Burrows-Wheeler transform (BWT) algorithm proposed by Crochemore et al. (Crochemore et al., 2015). Our extension is twofold: we first show how to compute simultaneously the longest common prefix (LCP) array as well as the BWT, using constant additional space; we then show how to build the LCP array directly in compressed representation using Elias coding, still using constant additional space and with no asymptotic slowdown. Furthermore, we provide a time/space tradeoff for our algorithm when additional memory is allowed. Our algorithm runs in quadratic time, as does Crochemore et al.'s, and is supported by interesting properties of the BWT and of the LCP array, contributing to our understanding of the time/space tradeoff curve for building indexing structures.

Keywords: Burrows-Wheeler transform, LCP array, In-place algorithms, Compressed LCP array, Elias coding

1. Introduction

There have been many articles [3, 4, 5, 7, 8] about building the Burrows-Wheeler transform (BWT) [9] and the longest common prefix (LCP) array. For

[☆] A preliminary version of this work appeared in IWOCA 2015 [2].

[*] Corresponding author
Email addresses: louza@ic.unicamp.br (Felipe A. Louza), travis.gagie@mail.udp.cl (Travis Gagie), gpt@ic.unicamp.br (Guilherme P. Telles)

example, Belazzougui [10] showed how we can compute the BWT and the (permuted) LCP array of a string T of length n over an alphabet of size σ in linear time and $O(n \log \sigma)$ bits of space (see also [11, 12]). Navarro and Nekrich [13] and Policriti, Gigante and Prezza [14] showed how to build the BWT in compressed space and, respectively, $O(n \log n / \log \log n)$ worst-case time and average-case time proportional to the length of the compressed representation of T.

The most space-efficient BWT construction algorithm currently known, however, is due to Crochemore *et al.* [1]: it builds the BWT in place — *i.e.*, replacing the input string with the BWT— in $O(n^2)$ time for unbounded alphabets using only a constant number of $\Omega(\log n)$ bit words of additional memory (*i.e.*, four integer variables and one character variable). Unlike most BWT-construction algorithms, this one is symmetric to the BWT inversion. Its simplicity and elegance make it very attractive from a theoretical point of view and it is interesting as one extreme of the time/space tradeoff curve for building BWTs. Because a quadratic time bound is impractical, however, Crochemore *et al.* showed how their algorithm can be speeded up at the cost of using more space.

Closely related to the BWT, the suffix array (SA) [15, 16] may be constructed by many algorithms in linear time (see [17, 18, 19] for reviews). Franceschini and Muthukrishnan [20] presented a suffix array construction algorithm that runs in $O(n \log n)$ time using constant additional space. The LCP array can be computed in linear time together with SA during the suffix sorting [21, 22] or independently given T and SA as input [5, 6, 7] or given the BWT [23, 8]. Table 1 summarizes the most closely related algorithms' bounds.[1]

In this article we show how Crochemore *et al.*'s algorithm can be extended to compute also the longest common prefix (LCP) array of a string T of length n. Specifically, we show how, given $\mathsf{BWT}(T[i+1, n-1])$ and $\mathsf{LCP}(T[i+1, n-1])$ and

[1] Although the authors did not mention it, it seems likely Navarro and Nekrich's and Policriti *et al.*'s algorithms can also be made to reuse the space occupied by the text for the BWT. With that modification, their $nH_k(T) + o(n \log \sigma)$ space bounds in Table 1 can be made $o(n \log \sigma)$.

Table 1: Summary of related works and their theoretical bounds. The last column shows the additional space used on top of what is needed to store the input and the output. Belazzougui's algorithm [10] was randomized but has been made deterministic [11, 12]. Navarro and Nekrich's algorithm [13] uses $nH_0(T) + o(n \log \sigma)$ bits on top of the text, where $H_k(T) \leq \lg \sigma$ is the kth-order empirical entropy of T. Policriti *et al.*'s algorithm [14] uses $nH_k(T) + n + O(\sigma \log n) + o(n \log \sigma)$ bits on top of the text and runs in $O(n(H_k(T) + 1))$ time in the average case. For discussion of empirical entropy, see, e.g., [24, 25, 26]. For simplicity, in this table we assume $\sigma \in \omega(1) \cap o(n/\log n)$.

	BWT	LCP	SA	time	additional space
Belazzougui [10]	✓	✓		$O(n)$	$O(n \log \sigma)$ bits
Navarro and Nekrich [13]	✓			$O(n \log n / \log \log n)$	$nH_k(S) + o(n \log \sigma)$ bits
Policriti *et al.* [14]	✓			$O(n(H_k(S) + 1)(\log n / \log \log n)^2)$	$nH_k(S) + o(n \log \sigma)$ bits
Crochemore *et al.* [1]	✓			$O(n^2)$	$O(1)$
Franceschini and Muthukrishnan [20]			✓	$O(n \log n)$	$O(1)$
Fischer [21]		✓	✓	$O(n)$	$O(n \log n)$ bits
Louza *et al.* [22]		✓	✓	$O(n\sigma)$	$O(\sigma \log n)$ bits
Our algorithm	✓	✓		$O(n^2)$	$O(1)$

$T[i]$, we can compute $\mathsf{BWT}(T[i, n-1])$ and $\mathsf{LCP}(T[i, n-1])$ using $O(n-i)$ time and constant extra space on top of what is needed to store $\mathsf{BWT}(T[i, n-1])$ and $\mathsf{LCP}(T[i, n-1])$. Our construction algorithm has many of the nice properties of Crochemore *et al.*'s original: it is conceptually simple and in-place, it allows practical time-space tradeoffs [2], and we can compute some compressed encodings of $\mathsf{LCP}(T)$ directly. This is particularly interesting because in practice the LCP array can be compressed by nearly a logarithmic factor. Computing the BWT and LCP together in small space is interesting, for example, when building compressed suffix trees (see, e.g. [27, 28, 29, 30, 31, 32]), which are space-efficient versions of the classic linear-space suffix tree [33] that is often based on the BWT and LCP.

[2] We are aware that the LCP array and the BWT array can be computed with similar worst-case bounds by using a combination of Franceschini and Muthukrishnan's algorithm to build the SA, then computing the LCP naively in $O(n^2)$ time overwriting the SA, and finally using Crochemore *et al.*'s algorithm to compute the BWT overwriting the text. We still think our algorithm is interesting, however, because of its simplicity — the C implementation fits in a single page — and its offer of encoding and tradeoffs.

There exist external memory algorithms that compute the BWT [45, 46] and the LCP array [41, 42, 43, 44]. In particular, Bauer *et al.* [47] and Cox *et al.* [48] showed how to construct the BWT and the LCP array simultaneously for string collections. They compute the LCP values and process the BWT in a order similar to the one we use for the algorithm in this article, but their solution uses auxiliary memory and partitions the output into buckets to address external-memory access issues. Tischler [49] introduced an external-memory algorithm that computes the Elias γ-coded [50] permuted LCP given the BWT and the sampled inverse suffix array as input. For further discussion, we refer the reader to recent books by Ohlebusch [24], Mäkinen *et al.* [25] and Navarro [26].

The rest of the article is organized as follows. In Section 2 we introduce concepts and notations. In Section 3 we review the in-place BWT algorithm by Crochemore *et al.*. In Section 4 we present our algorithm and in Section 5 we show how the LCP can be constructed in compressed representation. In Section 6 we provide a tradeoff between time and space for our algorithm when additional memory is allowed. In Section 7 we conclude the article and we leave an open question.

2. Background

Let Σ be an ordered alphabet of σ symbols. We denote the set of every nonempty string of symbols in Σ by Σ^+. We use the symbol $<$ for the lexicographic order relation between strings. Let \$ be a symbol not in Σ that precedes every symbol in Σ. We define $\Sigma^{\$} = \{T\$ \mid T \in \Sigma^+\}$.

The i-th symbol in a string T will be denoted by $T[i]$. Let $T = T[0]T[1]\dots T[n-1]$ be a string of length $|T| = n$. A substring of T will be denoted by $T[i,j] = T[i]\dots T[j]$, $0 \leq i \leq j < n$. A prefix of T is a substring of the form $T[0,k]$ and a suffix is a substring of the form $T[k,n-1]$, $0 \leq k < n$. The suffix $T[k,n-1]$ will be denoted by T_k.

Suffix array, LCP array and the BWT

A suffix array for a string provides the lexicographic order for all its suffixes. Formally, a suffix array SA for a string $T \in \Sigma^{\$}$ of size n is an array of integers $\mathsf{SA} = [i_0, i_1, \ldots, i_{n-1}]$ such that $T_{i_0} < T_{i_1} < \ldots < T_{i_{n-1}}$ [15, 16].

Let $lcp(S, T)$ be the length of the longest common prefix of two strings S and T in $\Sigma^{\$}$. The LCP array for T stores the value of lcp for suffixes pointed by consecutive positions of a suffix array. We define $\mathsf{LCP}[0] = 0$ and $\mathsf{LCP}[i] = lcp(T_{\mathsf{SA}[i]}, T_{\mathsf{SA}[i-1]})$ for $1 \leq i < n$.

The BWT of a string T can be constructed by listing all the n circular shifts of T, lexicographically sorting them, aligning the shifts columnwise and taking the last column [9]. The BWT is reversible and tends to group identical symbols in runs. It may also be defined in terms of the suffix array, to which it is closely related. Let the BWT of a string T be denoted simply by BWT. We define $\mathsf{BWT}[i] = T[\mathsf{SA}[i] - 1]$ if $\mathsf{SA}[i] \neq 0$ or $\mathsf{BWT}[i] = \$$ otherwise.

The first column of the conceptual matrix of the BWT will be referred to as F, and the last column will be referred to as L. The LF-mapping property of the BWT states that the i^{th} occurrence of a symbol $\alpha \in \Sigma$ in L corresponds to the i^{th} occurrence of α in F.

Some other relations between the SA and the BWT are the following. It is easy to see that $L[i] = \mathsf{BWT}[i]$ and $F[i] = T[\mathsf{SA}[i]]$. Moreover, if the first symbol of $T_{\mathsf{SA}[i]}$, $T[\mathsf{SA}[i]] = \alpha$, is the k^{th} occurrence of α in F, then j is the position of $T_{\mathsf{SA}[i]+1}$ in SA (*i.e.* j is the rank of $T_{\mathsf{SA}[i]+1}$) such that $L[j]$ corresponds to the k^{th} occurrence of α in L.

As an example, Figure 1 shows the circular shifts, the sorted circular shifts, the SA, the LCP, the BWT and the sorted suffixes for $T = \text{BANANA\$}$.

The range minimum query (rmq) with respect to the LCP is the smallest lcp value in an interval of a suffix array. We define $rmq(i, j) = \min_{i<k\leq j}\{\mathsf{LCP}[k]\}$, for $0 \leq i < j < n$. Given a string T of length n and its LCP array, it is easy to see that $lcp(T_{\mathsf{SA}[i]}, T_{\mathsf{SA}[j]}) = rmq(i, j)$.

	circular shifts	sorted circular shifts					sorted suffixes
i		F	L	SA	LCP	BWT	$T_{\mathsf{SA}[i]}$
0	BANANA\$	\$BANANA		6	0	A	\$
1	\$BANANA	A\$BANAN		5	0	N	A\$
2	A\$BANAN	ANA\$BAN		3	1	N	ANA\$
3	NA\$BANA	ANANA\$B		1	3	B	ANANA\$
4	ANA\$BAN	BANANA\$		0	0	\$	BANANA\$
5	NANA\$BA	NA\$BANA		4	0	A	NA\$
6	ANANA\$B	NANA\$BA		2	2	A	NANA\$

Figure 1: SA, LCP and BWT for $T =$ BANANA\$.

Elias coding

The Elias γ-code of a positive number $\ell \geq 1$ is composed of the unary-code of $\lfloor \log_2 \ell \rfloor + 1$ (a sequence of $\lfloor \log_2 \ell \rfloor$ 0-bits ended by one 1-bit), followed by the binary code of ℓ without the most significant bit [51]. The γ-code encodes ℓ in $2\lfloor \log_2 \ell \rfloor + 1$ bits. For instance, $\gamma(4) = 00100$, since the unary code for $\lfloor \log_2 4 \rfloor + 1 = 3$ is 001 and 4 in binary is 100.

The Elias δ-code of ℓ is composed of the γ-code of $1 + \lfloor \log_2 \ell \rfloor$, followed by the binary code of ℓ without the most significant bit. The δ-coding represents ℓ using $2\lfloor \log_2(\lfloor \log_2 \ell \rfloor + 1) \rfloor + 1 + \lfloor \log_2 \ell \rfloor$ bits, which is asymptotically optimal [50]. For instance, $\delta(9) = 00100001$, since $\gamma(\lfloor \log_2 9 \rfloor + 1) = 00100$ and 9 in binary is 1001.

Decoding a γ-encoded number ℓ_γ requires finding the leftmost 1-bit in the unary code of $\lfloor \log_2 \ell \rfloor + 1$, and interpreting the next $\ell - 1$ bits as a binary code. Decoding a δ-encoded number ℓ_δ requires decoding a γ-code and then reading the proper number of following bits as a binary code. Both decodings may be performed in constant time in a CPU having instructions for counting the number of leading zeros and shifting a word by an arbitrary number of bits.

3. In-place BWT

The algorithm by Crochemore *et al.* [1] overwrites the input string T with the BWT as it proceeds by induction on the suffix length.

Let $\mathsf{BWT}(T_s)$ be the BWT of the suffix T_s, stored in $T[s, n-1]$. The base cases are the two rightmost suffixes, for which $\mathsf{BWT}(T_{n-2}) = T_{n-2}$ and $\mathsf{BWT}(T_{n-1}) = T_{n-1}$. For the inductive step, the authors have shown that the position of \$ in $\mathsf{BWT}(T_{s+1})$ is related to the rank of T_{s+1} among the suffixes $T_{s+1},\ldots,T_{n-1}$ (local rank), thus allowing for the construction of $\mathsf{BWT}(T_s)$ even after $T[s+1, n-1]$ has been overwritten with $\mathsf{BWT}(T_{s+1})$. The algorithm comprises four steps.

1 Find the position p of \$ in $T[s+1, n-1]$. Evaluating $p - s$ gives the local rank of T_{s+1} that originally was starting at position $s+1$.

2 Find the local rank r of the suffix T_s using just symbol $c = T[s]$. To this end, sum the number of symbols in $T[s+1, n-1]$ that are strictly smaller than c with the number of occurrences of c in $T[s+1, p]$ and with s.

3 Store c into $T[p]$, replacing \$.

4 Shift $T[s+1, r]$ one position to the left. Write \$ in $T[r]$.

The algorithm runs in $O(n^2)$ time using constant space memory. Furthermore, the algorithm is also in-place since it uses $O(1)$ additional memory and overwrites the input text with the output BWT.

4. LCP array in constant space

Our algorithm computes both the BWT and the LCP array by induction on the length of the suffix. The BWT construction is the same as proposed by Crochemore *et al.* [1]. Let us first introduce an overview of our algorithm.

At a glance, the LCP evaluation works as follows. Suppose that $\mathsf{BWT}(T_{s+1})$ and the LCP array for the suffixes $\{T_{s+1}, \ldots, T_{n-1}\}$, denoted by $\mathsf{LCP}(T_{s+1})$,

have already been built. Adding the suffix T_s to the solution requires evaluating exactly two values of lcp, involving the two suffixes that will be adjacent to T_s.

The first lcp value involves T_s and the largest suffix T_a in $\{T_{s+1}, \ldots, T_{n-1}\}$ that is smaller than T_s. Fortunately, $\mathsf{BWT}(T_{s+1})$ and $\mathsf{LCP}(T_{s+1})$ are sufficient to compute such value. Recall that if the first symbol of T_a is not equal to the first symbol of T_s then $lcp(T_a, T_s) = 0$. Otherwise $lcp(T_a, T_s) = lcp(T_{a+1}, T_{s+1}) + 1$ and the rmq may be used, since both T_{a+1} and T_{s+1} are already in $\mathsf{BWT}(T_{s+1})$. We know that the position of T_{s+1} is p from Step 1 of the in-place BWT in Section 3. Then it is enough to find, in $\mathsf{BWT}(T_{s+1})$, the position of T_{a+1}, which stores the symbol corresponding to the first symbol of T_a.

The second lcp value involves T_s and the smallest suffix T_b in $\{T_{s+1}, \ldots, T_{n-1}\}$ that is larger than T_s. It may be computed in a similar fashion.

Basic algorithm

Suppose that $\mathsf{BWT}(T_{s+1})$ and $\mathsf{LCP}(T_{s+1})$ have already been built and are stored in $T[s+1, n-1]$ and $\mathsf{LCP}[s+1, n-1]$, respectively. Adding T_s, whose rank is r, to the solution requires updating $\mathsf{LCP}(T_{s+1})$: by first shifting $\mathsf{LCP}[s+1, r]$ one position to the left and then computing the new values of $\mathsf{LCP}[r]$ and $\mathsf{LCP}[r+1]$, which refer to the two suffixes adjacent to T_s in $\mathsf{LCP}(T_s)$.

The value of $\mathsf{LCP}[r]$ is equal to the lcp of T_s and T_a in $\mathsf{BWT}(T_{s+1})$. The rank of T_a is r and will be $r-1$ in $\mathsf{BWT}(T_s)$ after shifting. If the first symbol of T_a is equal to $T[s]$ then $\mathsf{LCP}[r] = lcp(T_{a+1}, T_{s+1}) + 1$, otherwise $\mathsf{LCP}[r] = 0$.

We can evaluate $lcp(T_{a+1}, T_{s+1})$ by the rmq function from the position of T_{a+1} to the position of T_{s+1}. We know that p is the position of T_{s+1} in $\mathsf{BWT}(T_{s+1})$. Then we must find the position p_{a+1} of T_{a+1} in $\mathsf{BWT}(T_{s+1})$.

Note that $T[p_{a+1}]$ corresponds to the first symbol of T_a. If $T[p_{a+1}] \neq T[s]$ then $lcp(T_a, T_s) = 0$, otherwise the value of $lcp(T_a, T_s)$ may be evaluated as $lcp(T_{a+1}, T_{s+1}) + 1 = rmq(p_{a+1}, p) + 1$.

The value of $\mathsf{LCP}[r+1]$ may be evaluated in a similar fashion. Let T_b be the suffix with rank $r+1$ in $\mathsf{BWT}(T_{s+1})$ (its rank will still be $r+1$ in $\mathsf{BWT}(T_s)$). We

must find the position p_{b+1} of T_{b+1} in $\mathsf{BWT}(T_{s+1})$ and then if $T[s] = T[p_{b+1}]$ compute $\mathsf{LCP}[r+1] = lcp(T_s, T_b) = lcp(T_{s+1}, T_{b+1}) + 1 = rmq(p, p_{b+1}) + 1$.

The algorithm proceeds by induction on the length of the suffix. It is easy to see that for the suffixes with length 1 and 2, the values in LCP will be always equal to 0. Let the current suffix be T_s $(0 \leq s \leq n-3)$. Our algorithm has new Steps 2', 2" and 4', added just after Steps 2 and 4, respectively, of the in-place BWT algorithm as follows:

2' Find the position p_{a+1} of the suffix T_{a+1}, such that suffix T_a has rank r in $\mathsf{BWT}(T_{s+1})$, and compute:

$$\ell_a = \begin{cases} rmq(p_{a+1}, p) + 1 & \text{if } T[p_{a+1}] = T[s] \\ 0 & \text{otherwise.} \end{cases}$$

2" Find the position p_{b+1} of the suffix T_{b+1}, such that suffix T_b has rank $r+1$ in $\mathsf{BWT}(T_{s+1})$, and compute:

$$\ell_b = \begin{cases} rmq(p, p_{b+1}) + 1 & \text{if } T[s] = T[p_{b+1}] \\ 0 & \text{otherwise.} \end{cases}$$

4' Shift $\mathsf{LCP}[s+1, r]$ one position to the left, store ℓ_a in $\mathsf{LCP}[r]$ and if $r+1 < n$ then store ℓ_b in $\mathsf{LCP}[r+1]$.

Computing ℓ_a and ℓ_b

To find p_{a+1} and p_{b+1} and to compute ℓ_a and ℓ_b in Steps 2' and 2", we use the following properties.

Lemma 1. *Let T_s be the suffix to be inserted in $\mathsf{BWT}(T_{s+1})$ at position r. Let $T_a \in \{T_{s+1}, \ldots, T_{n-1}\}$ be the suffix whose rank is r in $\mathsf{BWT}(T_{s+1})$, and let p_{a+1} be the position of T_{a+1}. If $p_{a+1} \notin [s+1, p)$ then $T[p_{a+1}] \neq T[s]$.*

Proof. The local rank of T_a in $\mathsf{BWT}(T_{s+1})$ is $r - s$. We know that $T[p_{a+1}]$ corresponds to the first symbol of T_a, and it follows from LF-mapping that the local rank of $T[p_{a+1}]$ is $r - s$ in $\mathsf{BWT}(T_{s+1})$. Then $T[p_{a+1}]$ is smaller than or

equal to $T[s]$, since T_s also has local rank $r-s$. If $T[p_{a+1}]$ is smaller than $T[s]$, p_{a+1} must be in $[s+1, n)$. However, if $T[p_{a+1}] = T[s]$ then p_{a+1} must precede the position where $T[s]$ will be inserted, *i.e.* the position p of T_{s+1}, otherwise the local rank of T_s would be smaller than $r-s$. Then if $T[p_{a+1}] = T[s]$ it follows that $p_{a+1} \in [s+1, p)$. □

We can use Lemma 1 to verify whether $T[p_{a+1}] = T[s]$ by simply checking if there is a symbol in $T[s+1, p-1]$ equal to $T[s]$. If no such symbol is found, $\ell_a = 0$, otherwise we need to compute $rmq(p_{a+1}, p)$. Furthermore, if we have more than one symbol in $T[s+1, p-1]$ equal to $T[s]$, the symbol whose local rank is $r-s$ will be the last symbol found in $T[s+1, p-1]$, *i.e.* the largest symbol in $T[s+1, p-1]$ smaller than $T[s]$. Then, to find such symbol we can simply perform a backward scan in T from $p-1$ to $s+1$ until we find the first occurrence of $T[p_{a+1}] = T[s]$. One can see that we are able, simultaneously, to compute the minimum function for the *lcp* visited values, obtaining $rmq(p_{a+1}, p)$ as soon as we find $T[p_{a+1}] = T[s]$.

Lemma 2. *Let T_s be the suffix to be inserted in $\mathsf{BWT}(T_{s+1})$ at position r. Let $T_b \in \{T_{s+1}, \ldots, T_{n-1}\}$ be the suffix whose rank is $r+1$ in $\mathsf{BWT}(T_{s+1})$, and let p_{b+1} be the position of T_{b+1}. If $p_{b+1} \notin (p, n-1]$ then $T[s] \neq T[p_{b+1}]$.*

The proof of Lemma 2 is similar to the proof of Lemma 1 and will be omitted. It is important to remember, though, that T_b will still have rank $r+1$ in $\mathsf{BWT}(T_s)$ (after inserting T_s).

The procedure to find ℓ_b uses Lemma 2 and computes $lcp(T_{s+1}, T_{b+1})$ in a similar fashion. It scans T from $p+1$ to $n-1$ until it finds the first occurrence of $T[p_{b+1}] = T[s]$, computing the minimum function to solve the rmq if such symbol is found.

The C source code presented in Figure 4 implements the algorithm using eight integer variables apart from the $n \log_2 \sigma$ bits used to store T and compute the BWT, and the $n \log_2 n$ bits used to compute the LCP array. This code is also available at `https://github.com/felipelouza/bwt-lcp-in-place`.

```
1  void compute_bwt_lcp(unsigned char *T, int n, int *LCP){
2  int i, p, r=1, s, p_a1, p_b1, l_a, l_b;
3  LCP[n-1] = LCP[n-2] = 0; // base cases
4
5  for (s=n-3; s>=0; s--) {
6
7    /*steps 1 and 2*/
8    p=r+1;
9     for (i=s+1, r=0; T[i]!=END_MARKER; i++)
10      if(T[i]<=T[s]) r++;
11    for (; i<n; i++)
12      if (T[i]<T[s]) r++;
13
14    /*step 2'*/
15    p_a1=p+s-1;
16    l_a=LCP[p_a1+1];
17    while (T[p_a1]!=T[s]) // rmq function
18      if (LCP[p_a1--]<l_a)
19        l_a=LCP[p_a1+1];
20    if (p_a1==s) l_a=0;
21    else l_a++;
22
23    /*step 2''*/
24    p_b1=p+s+1;
25    l_b=LCP[p_b1];
26    while (T[p_b1]!=T[s] && p_b1<n) // rmq function
27      if (LCP[++p_b1]<l_b)
28        l_b=LCP[p_b1];
29    if (p_b1==n) l_b=0;
30    else l_b++;
31
32    /*steps 3 and 4*/
33    T[p+s]=T[s];
34    for (i=s; i<s+r; i++) {
35      T[i]=T[i+1];
36      LCP[i]=LCP[i+1];
37    }
38    T[s+r]=END_MARKER;
39
40    /*step 4'*/
41    LCP[s+r]=l_a;
42    if (s+r+1<n)  // If r+1 is not the last position
43      LCP[s+r+1]=l_b;
44   }
45 }
```

Figure 2: BWT and LCP array construction algorithm

Example

As an example, consider T = BANANA\$ and $s = 1$. Figures 3 and 4 illustrate Steps 2' and 4', respectively. The values in red in columns LCP and BWT were still not computed. Suppose that we have computed BWT(T_2) and LCP(T_2). We then have $p = 6$ (Step 1) and the rank $r = 4$ (Step 2).

	s	LCP	BWT	sorted suffixes
	0	-	B	BANANA\$
$s \rightarrow$	1	-	A	ANANA\$
	2	0	A	\$
	3	0	N	A\$
$r \rightarrow$	4	1	N	ANA\$
$p_{a+1} \rightarrow$	5	0	A	NA\$
$p \rightarrow$	6	2	\$	NANA\$

Figure 3: After Step 2": T = BANANA\$ and $s = 1$.

Step 2' finds the first symbol equal to $T[s]$ (A) in $T[s+1, p-1]$ at position $p_{a+1} = 5$. It represents T_{a+1} = NA\$. In this case, the value of ℓ_a is calculated during the scan of T from $p - 1 = 5$ to $s + 1 = 2$, *i.e.* $\ell_a = rmq(p_{a+1}, p) = rmq(5, 6) = 2$. Step 2" does not find any symbol equal to $T[s]$ (A) in $T[p+1, n-1]$. Thus we know that $T[s] \neq T[p_{b+1}]$ and $\ell_b = 0$.

	s	LCP	BWT	sorted suffixes
	0	-	B	BANANA\$
	1	0	A	\$
	2	0	N	A\$
	3	1	N	ANA\$
$r \rightarrow$	4	$\ell_a = 3$	\$	ANANA\$
	5	$\ell_b = 0$	A	NA\$
	6	2	A	NANA\$

Figure 4: After Step 4': T = BANANA\$ and $s = 1$.

Step 3 stores $T[s]$ (A) at position $T[p]$, $p = 6$. Step 4 shifts $T[s+1, r]$ one position to the left and inserts \$ at position $T[r]$, $r = 4$. The last step, 4', shifts $\mathsf{LCP}[s+1, r]$ one position to the left and sets $\mathsf{LCP}[4] = \ell_a = 3$ and $\mathsf{LCP}[4+1] = \ell_b = 0$.

Theorem 1. *Given a string T of length n, we can compute its* BWT *in-place and* LCP *array simultaneously in $O(n^2)$ time using $O(1)$ additional space.*

Proof. The cost added by Steps 2' and 2" were two $O(n)$ time scans over T_{s+1} to compute the values of ℓ_a and ℓ_b, whereas Step 4' shifts the LCP by the same amount that BWT is shifted. Therefore, the time complexity of our algorithm remains the same as the in-place BWT algorithm, that is, $O(n^2)$. As for the space usage, our new algorithm needs only four additional variables to store positions p_{a+1} and p_{b+1} and the values of ℓ_a and ℓ_b, thus using constant space only. □

5. LCP array in compressed representation

The LCP array can be represented using less than $n \log n$ bits. Some alternatives for encoding the LCP array store its values in text order [52, 53], building an array that is known as permuted LCP (PLCP) [7]. Some properties of the PLCP will allow for encoding the whole array achieving better compression rates. However, most applications will require the LCP array itself, and will convert the PLCP to the LCP array [31]. Other alternatives for encoding the LCP will preserve its elements' order [54, 55].

Recall that to compute the BWT and the LCP array in constant space only sequential scans are performed. Therefore, the values in the LCP array can be easily encoded and decoded during such scans using a universal code, such as Elias δ-codes [25], with no need to further adjust the algorithm. Our LCP array representation will encode its values in the same order, and will be generated directly.

The algorithm will build the BWT and a compressed LCP array that will be called LCP_2. LCP_2 will be treated as a sequence of bits from this point on.

The *lcp* values will be δ-encoded during the algorithm such that consecutive intervals $\mathsf{LCP}_2[b_i, e_i]$ encode $lcp(T_{\mathsf{SA}[i]}, T_{\mathsf{SA}[i-1]}) + 1$. We add 1 to guarantee that the values are always positive integers and can be encoded using δ-codes. We will assume that decoding subtracts this 1 added by the encoding operation.

Suppose that $\mathsf{BWT}(T_{s+1})$ and $\mathsf{LCP}_2(T_{s+1})$ have already been built such that every value in $\mathsf{LCP}_2(T_{s+1})$ is δ-encoded and stored in $\mathsf{LCP}_2[b_{s+1}, e_{n-1}]$. Adding T_s to the solution requires evaluating the values of ℓ_a and ℓ_b computed in Steps 2' and 2" and the length of the shift to be performed in $\mathsf{LCP}_2[b_{s+1}, e_{r-1}]$ by Step 4'.

Modified Step 2'

We know by Lemma 1 that if there is no symbol in $T[s+1, p-1]$ equal to $T[s]$, then $\ell_a = 0$, which is encoded as $\delta(0+1) = 1$. Otherwise, if $T[s]$ occurs at position $p_{a+1} \in [s+1, p)$, we may compute $rmq(p_{a+1}, p)$ as the minimum value encoded in $\mathsf{LCP}_2[b_{p_{a+1}+1}, e_p]$. We use two extra variables to store the positions b_{s+1} and e_p of LCP_2 corresponding to the beginning of the encoded $lcp(T_{\mathsf{SA}[s+1]}, T_{\mathsf{SA}[s]}) + 1$ and the ending of the encoded $lcp(T_{\mathsf{SA}[p]}, T_{\mathsf{SA}[p-1]}) + 1$. These two variables are easily updated at each iteration.

As our algorithm performs a backward scan in T to find $T[p_{a+1}]$, we cannot compute the *rmq* function decoding the *lcp* values during this scan. Therefore, we first search for position p_{a+1} scanning T. Then, if p_{a+1} exists, the first bit of $\mathsf{LCP}_2[p_{a+1}+1]$ is found by decoding and discarding the first $p_{a+1}-s+1$ encoded values from b_{s+1}. At this point $rmq(p_{a+1}, p)$ may be evaluated by finding the minimum encoded value from $\mathsf{LCP}_2[p_{a+1}+1]$ to $\mathsf{LCP}_2[e_p]$. At the end, we add 1 to obtain ℓ_a.

Modified Step 2"

The algorithm performs a forward scan in T to find the position $p_{b+1} \in (p, n-1]$. Analogously to Modified Step 2', we know by Lemma 2 that if $T[s]$ does not occur in $T[p+1, n-1]$ then $\ell_b = 0$, which is encoded as $\delta(0+1)$. Otherwise, the *rmq* over the encoded *lcp* values may be computed during this

scan. The value of $rmq(p, p_{b+1})$ is computed decoding the values in LCP_2 one by one, starting at position $e_p + 1$ and continuing up to position p_{b+1} in T. At the end, we add 1 to obtain ℓ_b.

Modified Step 4'

The amount of shift in the compressed $\mathsf{LCP}_2(T_{s+1})$ must account for the sizes of $\delta(\ell_a + 1)$, of $\delta(\ell_b + 1)$ and of the encoding of the *lcp* value in position b_r, which represents $\delta(lcp(T_{\mathsf{SA}[r]}, T_{\mathsf{SA}[r-1]}) + 1)$ and will be overwritten by ℓ_b. We use two auxiliary integer variables to store positions b_{r+1} and e_{r+1}. We compute b_{r+1} and e_{r+1} by scanning LCP_2 from b_{s+1} up to finding e_{r+1}, by counting the encoding lengths one by one. The values in $\mathsf{LCP}_2[b_{r+1}, e_{r+1}]$ are set to 0 and $\mathsf{LCP}_2[b_{s+1}, b_{r+1} - 1]$ is shifted $|\delta(\ell_a + 1)| + |\delta(\ell_b + 1)| - (e_{r+1} - b_{r+1})$ positions to the left. To finish, the values of $\delta(\ell_a + 1)$ and $\delta(\ell_b + 1)$ are inserted into their corresponding positions b_r and b_{r+1} in LCP_2.

Theorem 2. *Given a string T of length n, we can compute its* BWT *in-place and* LCP *array compressed in $O(n \log\log n)$ bits, in the average case, in $O(n^2)$ time using $O(1)$ additional space.*

Proof. The cost added by the modifications in Steps 2', 2" and 4' is constant since the encoding and decoding operations are performed in $O(1)$ time and the left-shifting of the encoded *lcp* values in Step 4' is done word-size. Therefore, the worst-case time complexity of the modified algorithm remains $O(n^2)$. As for the space usage, the expected value of each LCP array entry is $O(\log n)$ for random texts [56] and for more specific domains, such as genome sequences and natural language, this limit has been shown empirically [57]. Therefore, in the average case our LCP array representation uses $O(n \log\log n)$ bits, since we are using Elias δ-coding [50]. In the worst case, when the text has only the same symbols, the LCP array still requires $n \log n$ bits since $\sum_{i=0}^{n-1} \log(i) = \log(n!) = \Theta(n \log n)$. □

6. Tradeoff

Crochemore *et al.* showed how, given $k \leq n$, we can modify their algorithm to run in $O((n^2/k+n)\log k)$ time using $O(k\sigma_k)$ space, where σ_k is the maximum number of distinct characters in a substring of length k of the text. The key idea is to insert characters from the text into the BWT in batches of size k, thereby using $O(1)$ scans over the BWT for each batch, instead of for each character. Their algorithm can be modified further to output, for each batch of k characters, a list of the k positions where those characters should be inserted into the current BWT, and the position where the \$ should be afterward [58]. (This modification has not yet been implemented, so neither has the tradeoff we describe below.)

From the list for a batch, with $O(1)$ passes over the current BWT using $O(k\sigma_k)$ additional space, we can compute in $O((n+k)\log k)$ time the intervals in the current LCP array on which we should perform *rmq*s when inserting that batch of characters and updating the LCP array, and with $O(1)$ more passes in $O(n)$ time using $O(k)$ additional space, we can perform those *rmq*s. The only complication is that we may update the LCP array in the middle of one of those intervals, possibly reducing the result of future *rmq*s on it. This is easy to handle with $O(k)$ more additional space, however, and does not change our bounds. Analogous to Crochemore *et al.*'s tradeoff, therefore, we have the following theorem:

Theorem 3. *Given a string T of length n and $k \leq n$, we can compute its* BWT *in-place and* LCP *array simultaneously in $O((n^2/k+n)\log k)$ time using $O(k\sigma_k)$ additional space, where σ_k is the maximum number of distinct characters in a substring of length k of the text.*

7. Conclusion

We have shown how to compute the LCP array together with the BWT using constant space. Like its predecessor, our algorithm is quite simple and it

builds on interesting properties of the BWT and of the LCP array. Moreover, we show how to compute the LCP array directly in compressed representation with no asymptotic slowdown using Elias coding, and we provide a time/space tradeoff for our algorithm when additional memory is allowed. We note that our algorithm can easily construct the suffix array using constant space, with no overhead on the running time. We also note that very recently there has been exciting work on obtaining better bounds via randomization [59].

We leave as an open question whether our algorithm can be modified to compute simultaneously the BWT and the permuted LCP in compressed form, which takes only 2n + $o(n)$ bits, while using quadratic or better time and only $O(n)$ bits on top of the space that initially holds the string and eventually holds the BWT.

Acknowledgments

We thank the anonymous reviewers and Meg Gagie for comments that improved the presentation of the manuscript. We thank Djamal Belazzougui, Juha Kärkkäinen, Gonzalo Navarro, Yakov Nekrich and Nicola Prezza for helpful discussions.

Funding: F.A.L. acknowledges the financial support of CAPES and CNPq [grant number 162338/2015-5]. T.G. acknowledges the support of the Academy of Finland and Basal Funds FB0001, Conicyt, Chile. G.P.T. acknowledges the support of CNPq.

References

[1] M. Crochemore, R. Grossi, J. Kärkkäinen, G. M. Landau, Computing the Burrows-Wheeler transform in place and in small space, J. Discret. Algorithms 32 (2015) 44–52.

[2] F. A. Louza, G. P. Telles, Computing the BWT and the LCP array in constant space, in: Proc. IWOCA, 2015, pp. 312–320.

[3] D. Okanohara, K. Sadakane, A linear-time Burrows-Wheeler transform using induced sorting, in: Proc. SPIRE, 2009, pp. 90–101.

[4] G. Tischler, Faster average case low memory semi-external construction of the Burrows-Wheeler Transform., in: Proc. ICABD, 2014, pp. 61–68.

[5] T. Kasai, G. Lee, H. Arimura, S. Arikawa, K. Park, Linear-time longest-common-prefix computation in suffix arrays and its applications, in: Proc. CPM, 2001, pp. 181–192.

[6] G. Manzini, Two space saving tricks for linear time LCP array computation, in: Proc. SWAT, 2004, pp. 372–383.

[7] J. Kärkkäinen, G. Manzini, S. J. Puglisi, Permuted longest-common-prefix array, in: Proc. CPM, 2009, pp. 181–192.

[8] T. Beller, S. Gog, E. Ohlebusch, T. Schnattinger, Computing the longest common prefix array based on the Burrows-Wheeler transform, J. Discret. Algorithms 18 (2013) 22–31.

[9] M. Burrows, D. Wheeler, A block-sorting lossless data compression algorithm, Tech. rep., Digital SRC Research Report (1994).

[10] D. Belazzougui, Linear time construction of compressed text indices in compact space, in: Proc. STOC, 2014, pp. 184–193.

[11] D. Belazzougui, F. Cunial, J. Kärkkäinen, V. Mäkinen, Linear-time string indexing and analysis in small space, CoRR abs/1609.06378 (2016).

[12] J. I. Munro, G. Navarro, Y. Nekrich, Space-efficient construction of compressed indexes in deterministic linear time, to appear in Proc. SODA, 2017.

[13] G. Navarro, Y. Nekrich, Optimal dynamic sequence representations, SIAM J. Comput. 43 (5) (2014) 1781–1806.

[14] A. Policriti, N. Gigante, N. Prezza, Average linear time and compressed space construction of the burrows-wheeler transform, in: Proc. LATA, 2015, pp. 587–598.

[15] U. Manber, E. W. Myers, Suffix arrays: A new method for on-line string searches, SIAM J. Comput. 22 (5) (1993) 935–948.

[16] G. H. Gonnet, R. A. Baeza-Yates, T. Snider, New indices for text: Pat trees and pat arrays, in: Information Retrieval, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1992, pp. 66–82.

[17] S. J. Puglisi, W. F. Smyth, A. H. Turpin, A taxonomy of suffix array construction algorithms, ACM Comp. Surv. 39 (2) (2007) 1–31.

[18] J. Dhaliwal, S. J. Puglisi, A. Turpin., Trends in suffix sorting: A survey of low memory algorithms, in: Proc. ACSC, 2012, pp. 91–98.

[19] J. Kärkkäinen, Suffix array construction, in: Encyclopedia of Algorithms, Springer, 2016, pp. 2141–2144.

[20] G. Franceschini, S. Muthukrishnan, In-place suffix sorting, in: Proc. ICALP, 2007, pp. 533–545.

[21] J. Fischer, Inducing the LCP-Array, in: Proc. WADS, 2011, pp. 374–385.

[22] F. A. Louza, S. Gog, G. P. Telles, Optimal suffix sorting and LCP array construction for constant alphabets, Inf. Process. Lett. 118 (2017) 30–34.

[23] S. Gog, E. Ohlebusch, Fast and lightweight LCP-array construction algorithms, in: Proc. ALENEX, 2011, pp. 25–34.

[24] E. Ohlebusch, Bioinformatics Algorithms: Sequence Analysis, Genome Rearrangements, and Phylogenetic Reconstruction, Oldenbusch Verlag, 2013.

[25] V. Mäkinen, D. Belazzougui, F. Cunial, A. I. Tomescu, Genome-Scale Algorithm Design, Cambridge University Press, 2015.

[26] G. Navarro, Compact Data Structures: A Practical Approach, Cambridge University Press, 2016.

[27] K. Sadakane, Compressed suffix trees with full functionality, Theory Comput. Systems 41 (4) (2007) 589–607.

[28] E. Ohlebusch, J. Fischer, S. Gog, CST++, in: Proc. SPIRE, 2010, pp. 322–333.

[29] J. Fischer, Combined data structure for previous- and next-smaller-values, Theor. Comput. Sci. 412 (22) (2011) 2451–2456.

[30] A. Abeliuk, R. Cánovas, G. Navarro, Practical compressed suffix trees, Algorithms 6 (2) (2013) 319–351.

[31] S. Gog, E. Ohlebusch, Compressed suffix trees: Efficient computation and storage of LCP-values, ACM J. Experiment. Algorithmics 18.

[32] G. Navarro, L. Russo, Fast fully-compressed suffix trees, in: Proc. DCC, 2014, pp. 283–291.

[33] P. Weiner, Linear pattern matching algorithms, in: Proc. SWAT, 1973, pp. 1–11.

[34] A. Crauser, P. Ferragina, A theoretical and experimental study on the construction of suffix arrays in external memory, Algorithmica 32 (1) (2002) 1–35.

[35] R. Dementiev, J. Kärkkäinen, J. Mehnert, P. Sanders, Better external memory suffix array construction, ACM J. Experiment. Algorithmics 12 (2008) 3.4:1–3.4:24.

[36] J. Kärkkäinen, D. Kempa, Engineering a lightweight external memory suffix array construction algorithm., in: Proc. ICABD, 2014, pp. 53–60.

[37] G. Nong, W. H. Chan, S. Zhang, X. F. Guan, Suffix array construction in external memory using d-critical substrings, ACM Trans. Inf. Syst. 32 (2014) 1:1–1:15.

[38] G. Nong, W. H. Chan, S. Q. Hu, Y. Wu, Induced sorting suffixes in external memory, ACM Trans. Inf. Syst. 33 (3) (2015) 12:1–12:15.

[39] W. J. Liu, G. Nong, W. H. Chan, Y. Wu, Induced sorting suffixes in external memory with better design and less space, in: Proc. SPIRE, 2015, pp. 83–94.

[40] J. Kärkkäinen, D. Kempa, S. J. Puglisi, Parallel external memory suffix sorting, in: Proc. CPM, 2015, pp. 329–342.

[41] J. Kärkkäinen, D. Kempa, LCP array construction in external memory, ACM J. Experiment. Algorithmics 21 (2016) 1–22.

[42] J. Kärkkäinen, D. Kempa, Faster external memory LCP array construction, in: Proc. ESA, 2016, pp. 61:1–61:16.

[43] J. Kärkkäinen, D. Kempa, LCP array construction using O(sort(n)) (or Less) I/Os, in: Proc. SPIRE, 2016, pp. 204–217.

[44] T. Bingmann, J. Fischer, V. Osipov, Inducing suffix and LCP arrays in external memory, ACM J. Experiment. Algorithmics 21 (2) (2016) 2.3:1–2.3:27.

[45] P. Ferragina, T. Gagie, G. Manzini, Lightweight data indexing and compression in external memory, Algorithmica 63 (3) (2012) 707–730.

[46] M. J. Bauer, A. J. Cox, G. Rosone, Lightweight algorithms for constructing and inverting the BWT of string collections, Theor. Comput. Sci. 483 (2013) 134–148.

[47] M. J. Bauer, A. J. Cox, G. Rosone, M. Sciortino, Lightweight LCP construction for next-generation sequencing datasets, in: Proc. WABI, 2012, pp. 326–337.

[48] A. J. Cox, F. Garofalo, G. Rosone, M. Sciortino, Lightweight LCP construction for very large collections of strings, J. Discret. Algorithms 37 (2016) 17–33.

[49] G. Tischler, T. Sakai, Low space external memory construction of the succinct permuted longest common prefix array, in: Proc. SPIRE, 2016, pp. 178–190.

[50] P. Elias, Universal codeword sets and representations of the integers, IEEE Trans. on Information Theory 21 (2) (1975) 194–203.

[51] I. H. Witten, A. Moffat, T. C. Bell, Managing Gigabytes: Compressing and Indexing Documents and Images, Second Edition, Morgan Kaufmann, 1999.

[52] K. Sadakane, Succinct representations of lcp information and improvements in the compressed suffix arrays, in: Proc. SODA, 2002, pp. 225–232.

[53] J. Fischer, Wee LCP, Inf. Process. Lett. 110 (8-9) (2010) 317–320.

[54] M. I. Abouelhoda, S. Kurtz, E. Ohlebusch, Replacing suffix trees with enhanced suffix arrays, J. Discret. Algorithms 2 (1) (2004) 53–86.

[55] N. R. Brisaboa, S. Ladra, G. Navarro, DACs: Bringing direct access to variable-length codes, Inf. Process. Manage. 49 (1) (2013) 392–404.

[56] J. Fayolle, M. D. Ward, Analysis of the average depth in a suffix tree under a Markov model, in: Proc. DMTCS, 2005, pp. 95–104.

[57] M. Léonard, L. Mouchard, M. Salson, On the number of elements to reorder when updating a suffix array, J. Discret. Algorithms 11 (2012) 87–99.

[58] J. Kärkkäinen. Personal communication.

[59] N. Prezza, In-place longest common extensions, CoRR abs/1608.05100 (2016).